16

Probability and Decision Theory

Man tracht und Gott lacht.
Yiddish proverb

Doubt is disagreeable, but certainty is ridiculous.
Voltaire

16.1 Probability Spaces

We assume a *universe* or *sample space* X, and a set $\mathcal{E}$ of subsets $A, B, C, \ldots$ of X, called *events*. We assume $\mathcal{E}$ is closed under countable unions (if $A_1, A_2, \ldots$ are events, so is $\cup_{i=1}^{\infty} A_i$), finite intersections (if $A_1, \ldots, A_n$ are events, so is $\cap_{i=1}^{n} A_i$), and complementation (if A is an event so is the set of elements of X that are not in A, which we write A^c). If A and B are events, we interpret $A \cap B = AB$ as the event "A and B both occur," $A \cup B$ as the event "A or B occurs," and A^c as the event "A does not occur."

For instance, suppose we flip a coin twice, the outcome being HH (heads on both), HT (heads on first and tails on second), TH (tails on first and heads on second), and TT (tails on both). The sample space is then $X = \{HH, TH, HT, TT\}$. Some events are $\{HH, HT\}$ (the coin comes up heads on the first toss), $\{TT\}$ (the coin comes up tails twice), and $\{HH, HT, TH\}$ (the coin comes up heads at least once).

The *probability* of an event $A \in \mathcal{E}$ is a real number $\mathrm{P}[A]$ such that $0 \leq \mathrm{P}[A] \leq 1$. We assume that $\mathrm{P}[X] = 1$, which says that with probability one *some* outcome occurs, and we also assume that if $A = \cup_{i=1}^{\infty} A_i$ where $A_i \in \mathcal{E}$ and the $\{A_i\}$ are disjoint (i.e., $A_i \cap A_j = \emptyset$ for all $i \neq j$), then $\mathrm{P}[A] = \sum_{i=1}^{\infty} P[A_i]$, which says that probabilities are countably additive over disjoint unions.

VISIT…

16.2 Conditional Probability

If A and B are events, and if the probability $\mathrm{P}[B]$ that B occurs is strictly positive, we define the *conditional probability* of A given B, denoted $\mathrm{P}[A|B]$, by

$$\mathrm{P}[A|B] = \frac{\mathrm{P}[AB]}{\mathrm{P}[B]}.$$

We say $B_1, \ldots, B_n$ are a *partition* of event B if $\cup_i B_i = B$ and $B_i B_j = \emptyset$ for $i \neq j$. We have

(a) If A and B are events, $\mathrm{P}[B] > 0$, and B implies A (i.e., $A \subseteq B$) then $\mathrm{P}[A|B] = 1$.

(b) If A and B are contradictory (i.e., $AB = \emptyset$), then $\mathrm{P}[A|B] = 0$.

(c) If $A_1, \ldots, A_n$ are a partition of event A, then

$$\mathrm{P}[A|B] = \sum_{i=1}^{n} \mathrm{P}[A_i|B].$$

(d) If $B_1, \ldots, B_n$ are a partition of the sample space X then

$$\mathrm{P}[A] = \sum_{i=1}^{n} \mathrm{P}[A|B_i]\mathrm{P}[B_i].$$

16.3 Bayes' Rule

If $B_1, \ldots, B_n$ is a partition of the sample space and if $\mathrm{P}[A], \mathrm{P}[B_k] > 0$ then

$$\mathrm{P}[B_k|A] = \frac{\mathrm{P}[A|B_k]\mathrm{P}[B_k]}{\sum_{i=1}^{n} \mathrm{P}[A|B_i]\mathrm{P}[B_i]}.$$

To see this, note that the denominator on the right hand side is just $\mathrm{P}[A]$, and the numerator is just $\mathrm{P}[AB_k]$ by definition.

As a special case, suppose A and B are events with $\mathrm{P}[A], \mathrm{P}[B], \mathrm{P}[B^c] > 0$. Then we have

$$\mathrm{P}[B|A] = \frac{\mathrm{P}[A|B]\mathrm{P}[B]}{\mathrm{P}[A|B]\mathrm{P}[B] + \mathrm{P}[A|B^c](1 - \mathrm{P}[B])}$$

16.4 Drug Testing

Bayes' Rule is useful because often we know $\mathrm{P}[A|B]$, $\mathrm{P}[A|B^c]$ and $\mathrm{P}[B]$, and we want to find $\mathrm{P}[B|A]$. For example, suppose 5% of the population uses drugs, and there is a drug test that is 95% accurate: it tests positive on a drug user 95% of the time, and it test negative on a drug non-user 95% of the time. Show that if an individual tests positive, the probability of his being a drug user is 50%. Hint: Let A be the event "is a drug user," let 'Pos' be the event "tests positive," let 'Neg' be the event "tests negative," and apply Bayes' Rule.

16.5 A Bolt Factory

In a bolt factory, machines A, B, and C manufacture 25%, 35% and 40% of total output and have defective rates of 5%, 4%, and 2%, respectively. A bolt is chosen at random and found to be defective. What are the probabilities that it was manufactured by machines A, B, and C?

16.6 Color Blindness

Suppose 5.0% of men are color blind and 0.25% of women are color blind. A person is chosen at random and found to be color blind. What is the probability the person in male?

16.7 Urns

A collection of $n+1$ urns, numbered from 0 to n, each contains n balls. Urn k contains k red and $n-k$ white balls. An urn is chosen at random and n balls are randomly chosen from it, the ball being replaced each time before another is chosen. Suppose all n balls are found to be red. What is the probability the next ball chosen from the urn will be red? Show that when n is large, this probability is approximately $(n+1)/(n+2)$. The last step is not obvious.

16.8 The Monty Hall Game

You are a contestant in a game show. The host says, "Behind one of those three doors is a new automobile. Behind the other two doors are goats (silent—like the automobile). You may choose any door." You choose door 1. The game host then opens door 2 and shows you that

there is a goat behind it. He then says, "now would you like to change your guess?" We will show that the answer is "no" if the game show host randomly opened one of the two other doors, but "yes" if he simply opened a door he knew did not have a car behind it.

Suppose there are n doors ($n \geq 3$). Let A_k be the event "car is behind door k." Let B_k be the event "host chooses door k, contestant did not choose door k, and door k has a goat behind it." We write $\mathrm{P}[E]$ for the probability of the event E and $\mathrm{P}[E|F]$ for the conditional probability of event E, given event F. Clearly $\mathrm{P}[A_k] = 1/n$ for $k = 1, \ldots, n$. By Bayes' Rule,

$$\mathrm{P}[A_j|B_k] = \frac{\mathrm{P}[B_k|A_j]\mathrm{P}[A_j]}{\mathrm{P}[B_k]} = \frac{\mathrm{P}[B_k|A_j]}{n\mathrm{P}[B_k]} \qquad j, k = 1, \ldots, n \qquad (16.1)$$

provided $\mathrm{P}[B_k] > 0$.

Suppose, without loss of generality, that the contestant chooses door 1, so he wins if A_1. Thus $\mathrm{P}[win] = \mathrm{P}[A_1] = 1/n$.

Now suppose the host chooses *randomly* from all doors except door 1. He chooses door $k > 1$ with probability $1/(n-1)$, and this has a goat behind it with probability $(n-1)/n$, so $\mathrm{P}[B_k] = 1/n$ for $k > 1$. If A_1, door k certainly has a goat behind it, so $\mathrm{P}[B_k|A_1] = 1/(n-1)$ for $k > 1$. From equation (16.1) we get

$$\mathrm{P}[A_1|B_k] = \frac{1/(n-1)}{n(1/n)} = 1/(n-1), \qquad k > 1.$$

If A_j for $j > 1$, then for $k \neq 1, j$, the probability of B_k conditional on A_j is $1/(n-1)$, since the host does not know what door has the car behind it, hence can choose any one of the doors $2, \ldots, n$, and door k certainly has a goat behind it for $k \neq j$. From $\mathrm{P}[B_k|A_j] = 1/(n-1)$ for $j \neq k$, $j > 1$, and equation (16.1) we get

$$\mathrm{P}[A_j|B_k] = \frac{1/(n-1)}{n(1/n)} = 1/(n-1), \qquad j \neq k, \quad j > 1.$$

which is the same as $\mathrm{P}[A_1|B_k]$—the contestant does not gain from switching. This is the professors' result.

But suppose the host *never* picks the door with car behind it. Now $\mathrm{P}[B_k|A_1] = 1/(n-1) = \mathrm{P}[B_k]$ for $k > 1$, so from equation (16.1) we have

$$\mathrm{P}[A_1|B_k] = \frac{1/(n-1)}{n[1/(n-1)]} = 1/n, \qquad k > 1.$$

For $j > 1$, $k \neq j$, $\mathrm{P}[B_k|A_j] = 1/(n-2)$, since the host must pick randomly from all doors except 1 and j. Thus from equation (16.1) we get

$$\mathrm{P}[A_j|B_k] = \frac{1/(n-2)}{n[1/(n-1)]} = \frac{(n-1)}{n(n-2)}, \qquad j \neq k, \quad j > 1.$$

Since $(n-1)/n(n-2) > 1/n$, the contestant should switch.

Note that if $n = 3$, these probabilities reduce to 1/3 and 2/3 in the one case and 1/2 and 1/2 in the other.

Here is another proof for the three-door case. Let C_i be the event "car behind door i." Let the doors be numbered by: contestant chose door 1, host chose door 2. Let H be the event "host picks animal." Suppose first that the host chooses *randomly* between the two doors. Then we have $\mathrm{P}[C_1] = \mathrm{P}[C_2] = \mathrm{P}[C_3] = 1/3$ since the labelling of the doors tells us nothing about what is behind them, so it is equiprobable that the animal is behind any door. Also $\mathrm{P}[H] = 2/3$, the probability that an animal is behind a randomly chosen door. Moreover $\mathrm{P}[H|C_1] = 1$, since if the car is behind door 1, the host *must* pick a door with an animal behind it. By Bayes' Rule,

$$\mathrm{P}[C_1|H] = \frac{\mathrm{P}[H|C_1]\mathrm{P}[C_1]}{\mathrm{P}[H]} = \frac{1(1/3)}{2/3} = 1/2.$$

Thus $\mathrm{P}[C_2|H] = 1/2$, and there is no reason to change doors. Now suppose the host chooses a door he *knows* has an animal behind it. Then $\mathrm{P}[C_2] = 0$, so $\mathrm{P}[C_3] = 2/3$, and $\mathrm{P}[H] = 1$. Then $\mathrm{P}[H|C_1] = 1$, and Bayes' Rule gives

$$\mathrm{P}[C_1|H] = \frac{\mathrm{P}[H|C_1]\mathrm{P}[C_1]}{\mathrm{P}[H]} = \frac{1(1/3)}{1} = 1/3.$$

Thus $\mathrm{P}[C_3|H] = 2/3$, and the contestant doubles the chances of winning by switching doors.

16.9 The Logic of Murder and Abuse

Let A be the event that Nicole had been repeatedly beaten by her husband ('abused' for short), let B be the event "was murdered," and let C be the event "was murdered by her husband." Suppose we know the following facts: (a) 5% of women are abused by their husbands; (b) 0.5% of women are murdered; (c) 0.025% of women are murdered by their

husbands; (d) 90% of women who are murdered by their husbands had been abused by their husbands; (e) a woman who is not murdered by her husband (whether or not she is murdered) is neither more nor less likely to have been abused by her husband than a randomly selected woman.

Nicole is found murdered, and it is ascertained that she had been abused by her husband. The defense attorneys for her husband show that the probability that a man who abuses his wife actually kills his wife is only 4.31%, so there is a strong presumption of innocence for him. The defense attorneys shows that there is in fact a 94.74% chance he murdered his wife, independent from any evidence other than he abused her. Here is how the two attorneys argued (the jury was well versed in probability theory, so they had no problem understanding the reasoning).

First, from Bayes' Rule,

$$\mathrm{P}[C|A] = \frac{\mathrm{P}[A|C]\mathrm{P}[C]}{\mathrm{P}[A|C]\mathrm{P}[C] + \mathrm{P}[A|C^c](1-\mathrm{P}[C])}.$$

This is the probability a man murders his wife if he has abused her. But from (d) above $\mathrm{P}[A|C] = 9/10$, from (c) $\mathrm{P}[C] = 1/400$, from (e) and (a) $\mathrm{P}[A|C^c] = \mathrm{P}[A] = 1/20$, so we find $\mathrm{P}[C|A] = 6/139 \approx 4.31\%$.

"I Object!" says the prosecutor. "The defense attorney ignores the fact that Nicole was *murdered*. What we *really* must know is $\mathrm{P}[C|AB]$, the probability a *murdered* woman who was abused by her husband was murdered by him." "But," splutters the astounded judge, "how could you calculate such a complex probability?" A large overhead projector is brought into the court, and the prosecuting attorney proceeds with the following calculation, the astute jurors taking mental notes. "We have," says the prosecutor,

$$\begin{aligned}\mathrm{P}[C|AB] &= \frac{\mathrm{P}[ABC]}{\mathrm{P}[AB]} = \frac{\mathrm{P}[A|BC]\mathrm{P}[BC]}{\mathrm{P}[ABC] + \mathrm{P}[ABC^c]} \\ &= \frac{\mathrm{P}[A|BC]\mathrm{P}[C|B]\mathrm{P}[B]}{\mathrm{P}[A|BC]\mathrm{P}[C|B]\mathrm{P}[B] + \mathrm{P}[A|BC^c]\mathrm{P}[C^c|B]\mathrm{P}[B]} \\ &= \frac{\mathrm{P}[A|C]\mathrm{P}[C|B]}{\mathrm{P}[A|C]\mathrm{P}[C|B] + \mathrm{P}[A](1-\mathrm{P}[C|B])},\end{aligned}$$

where $\mathrm{P}[A|BC^c] = \mathrm{P}[A]$ by (e). From (b) $\mathrm{P}[B] = 1/200$, so $\mathrm{P}[C|B] = \mathrm{P}[C]/\mathrm{P}[B] = 1/2$, so we have $\mathrm{P}[C|AB] = 18/19 = 94.74\%$

16.9.1 The Principle of Insufficient Reason

The *Principle of Insufficient Reason* says that if you are "completely ignorant" as to which among the states $A_1, \ldots, A_n$ will occur, then you should assign probability $1/n$ to each of the states. The argument in favor of the principle is strong (see Savage (1954) and Sinn (1980) for discussions), but there are some interesting arguments against it. For instance, suppose A_1 itself consists of m mutually exclusive events $A_{11}, \ldots, A_{1m}$. If you are "completely ignorant" concerning which of these occurs, then if $\mathrm{P}[A_1] = 1/n$, we should set $\mathrm{P}[A_{1i}] = 1/mn$. But are we not "completely ignorant" concerning which of

$$A_{11}, \ldots, A_{1m}, A_2, \ldots, A_n$$

occurs? If so, we should set each of these probabilities to $1/(n+m-1)$. If not, in what sense were we "completely ignorant" concerning the original states $A_1, \ldots, A_n$?

We will use the Principle of Insufficient Reason whenever it seems justified. As for the philosophical criticisms, we might do well to heed Spinoza's observation that "the philosophers kick up the dust and then complain that they cannot see." Or was it Leibniz...?

16.10 Ah, Those Kids

You see an acquaintance strolling down the street with a boy, whom he introduces as his son. You know he has another child, and he likes strolling with them equally. What is the probability the other child is a boy?

You see an acquaintance at a father and son dinner for 6^{th} graders. He introduces the boy as his son. You know he has another child, in the 8^{th} grade. What is the probability the other child is a boy?

16.11 The Greens and the Blacks

The Greens and the Blacks are playing bridge. After a deal, Mr. Brown, an onlooker, asks Mrs. Black: "Do you have an ace in your hand?" She nods yes. After the next deal, he asks her: "Do you have the ace of spades?" She nods yes again. In which of the two situations is Mrs. Black more likely to have at least one other ace in her hand? Calculate the exact probabilities in the two cases.

16.12 Laplace's Law of Succession

An urn contains a large number n of white and black balls, where the number of white balls is uniformly distributed between 0 and n. Suppose you pick out m balls with replacement, and r are white. Show that the probability of picking a white ball on the next draw is approximately $(r+1)/(m+2)$.

16.13 The Brain and Kidney Problem

A mad scientist is showing you around his foul-smelling laboratory. He motions to an opaque, formalin-filled jar. "This jar contains either a brain or a kidney," he exclaims. Searching around his workbench, he find a brain and adds it to the jar. He then picks one blob randomly from the jar, and it is a brain. What is the probability the remaining blob is a brain?

16.14 Sexual Harassment on the Job

Suppose you get a new supervisor every year on the job, and the probability that a new supervisor will harass you is 5% (i.e., one boss out of twenty is a slime-bag). Suppose your accumulated promotion and retirement benefits are $1000 per year of service, but you receive these only if you stay at the firm for 30 years. Suppose you would be willing to pay a maximum of $10,000 to avoid being harassed. What are the odds you quit your job to avoid being harassed? What are the expected number of years on the job you expect to be harassed?

16.15 The Value of Eye-Witness Testimony

A town has 100 taxis, 85 green taxis owned by the Green Cab Company and 15 blue taxies owned by the Blue Cab Company. On March 1, 1990 John Doe was struck by a speeding cab, and the only witness testified that the cab was blue rather than green. John Doe sued the Blue Cab Company. The judge instructed the jury and the lawyers at the start of the case that the reliability of a witness must be assumed to be 80% in a case of this sort, and that liability requires that the "preponderance of the evidence," meaning at least a 50% percent probability, be on the side of the plaintiff.

The lawyer for John Doe argued that Blue Cab Company should pay, since the witness's testimonial give a probability of 80% that Doe was struck by a blue taxi. The lawyer for the Blue Cab Company argued as follows. If the witness were shown all the cabs in town, he would have incorrectly identified 20% of the 85 green taxis (i.e., 17 of them) as blue, and correctly identified 80% of the 15 blue taxis (i.e., 12 of them) as blue. Thus of the 29 identifications of a taxi as blue, only 12 would be correct and 17 would be incorrect. Thus the preponderance of the evidence is in favor of the defendant–most likely Mr. Doe was hit by a green taxi.

Formulate the second lawyer's argument rigorously in terms of Bayes' Rule. Which argument do you think is correct, and if neither is correct, what is a good argument in this case?

16.16 The End of the World

Let A be the event "When the human race becomes extinct, fewer than 1 trillion (one million billion) persons will have lived." Let B be the event "Currently fewer than 100 billion persons have lived." Paleontological evidence tells us that $\mathrm{P}[B] \approx 1$, and quite clearly it is reasonable to suppose that $\mathrm{P}[A] \leq \mathrm{P}[B]$. Thus from Bayes' Rule we conclude that $\mathrm{P}[A|B] \leq \mathrm{P}[B|A]$ (prove it!). But also $\mathrm{P}[B|A] < \mathrm{P}[B|A'] < 10\%$, where A' is the event "When the human race becomes extinct, exactly 1 trillion persons will have lived," since if the current period is randomly chosen from among all the years that human beings are not extinct, the current period will be in the first 10% with probability 10%. This proves that $\mathrm{P}[A|B]$, which is approximately the probability that there will be fewer than a trillion people ever existing (since B is near certainty), is less that 10%. It follows that we should limit population growth, or humanity will become extinct very soon.

Now each of these steps appears valid, but the reasoning as a whole is quite absurd, at least to the minds of many. The renown physicist Freeman Dyson, for instance, criticizes it bitterly in (*Nature* 380 (1996):296). What do you think is wrong with the argument, if anything?

16.17 Bill and Harry

Bill and Harry each tell the truth 1/3 of the time. Bill says something, and Harry says "Bill just told the truth." What is the probability that Bill told the truth?

16.18 When Weakness is Strength

Many people have criticized the Darwinian notion of "survival of the fittest" by declaring that the whole thing is a simple tautology: whatever survives is ' fit' by definition! Defenders of the notion reply by noting that we can measure fitness (e.g., speed, strength, resistance to disease, aerodynamic stability) independent of survivability, so it becomes an empirical proposition that the fit survive. Indeed, under some conditions it may be simply false, as the game theorist Martin Shubik (1954) showed in the following ingenious problem.

Three marksmen, a, b, and c, are having a shootout. They draw straws to see who goes first, second and third, and then each shoots his gun, taking turns repeatedly in that order, until only one of them is left standing. The conditions are as follows. If a marksman is hit, he is eliminated; a is a perfect shot, b has 80% accuracy, and c has 50% accuracy. For notational purposes, let $\pi_i(xyz)$ be the probability of survival for player $i = a, b, c$ when all three are still in and they are in order xyz; for instance, $\pi_a(abc)$ is the probability a wins when the order is abc. Similarly, if only two remain, let $\pi_i(xy)$ be the probability of survival for player $i = x, y$ when only x and y remain, and it is x's turn to shoot.

(a) Show that when it is c's turn to shoot, he should shoot in the air rather than aiming at either of the other marksmen;

(b) Show that c, the least fit in terms of accuracy, is the most likely to survive, and find the probability that each of the marksmen survives;

(c) You may think that this counterintuitive result has something to do with the fact that the marksmen must shoot in a specific order. Show that if the marksman who fires is picked randomly in each round (i.e., the marksmen shoot at will, but not simultaneously), then the survivability of the marksmen is perfectly inverse to their accuracy.

16.19 Markov Chains

A *finite Markov chain* is a model that in each time period $t = 0, 1, \ldots$ can be any one of n *states*, such that if the system is in state i in one time period, there is a probability p_{ij} that the system is in state j in

the next time period. Thus for each i, we must have $\sum_j p_{ij} = 1$, since the system must go somewhere in each period. We call the $n \times n$ matrix $P = \{p_{ij}\}$ the *transition probability matrix* of the Markov chain, and each p_{ij} is called a *transition probability*. A *denumerable Markov chain* has an infinite number of states $t = 1, 2, \ldots$, and is otherwise the same. If we do not care whether the finite or denumerable case obtains, we speak simply of a *Markov chain*.

Many games can be viewed as Markov chains. Here are some examples:

(a) Suppose two gamblers have wealth k_1 and k_2 dollars, respectively, and in each period they play a game in which each has an equal chance of winning one dollar. The game continues until one player has no more wealth. Here the state of the system is the wealth w of player 1, and $p_{w,w+1} = p_{w,w-1} = 1/2$ for $0 < w < k_2$ and all other transition probabilities are zero.

(b) Suppose n agents play an game in which they are randomly paired in each period, and the stage game is a prisoner's dilemma. Players can remember the last k moves of their various partners. Players are also given one of r strategies, which determine their next move, depending on their current histories. When a player dies, which occurs with a certain probability, it is replaced by a new player that is a clone of a successful player. We can consider this a Markov chain in which the state of the system is the history, strategy, and score of each player, and the transition probabilities are just the probabilities of moving from one such state to another, given the players' strategies (§10.3).

(c) Suppose n agents play an game in which they are randomly paired in each period to trade. Each agent has an inventory of goods to trade, and a strategy indicating which goods the agent is willing to trade for which other goods. After trading, agents consume some of their inventory and produce more goods for their inventory, according to some consumption and production strategy. When an agent dies, it is replaced by a new agent with the same strategy and an empty inventory. If there is a maximum size inventory, and all goods are indivisible, we can consider this a finite Markov chain in which the state of the system is the strategy and inventory of each player and the transition probabilities are determined accordingly (§10.2).

(d) In a population of beetles, females have k offspring in each period with probability f_k and beetles live for n periods. The state of the

system is the fraction of males and females of each age. This is a denumerable Markov chain, where the transition probabilities are calculated from the birth and death rates of the beetles.

We are interested in the long run behavior of Markov chains. In particular, we are interested in the behavior of systems that we expect to attain a long run equilibrium of some type independent from its initial conditions. If such an equilibrium exists, we say the Markov chain is *ergodic*. In an ergodic system, history does not matter—every initial condition leads to the same long-run behavior. Nonergodic systems are history-dependent. It is intuitively reasonable that the repeated prisoner's dilemma and the trading model described above are ergotic. The gambler model is not ergodic, because the system could end up with either player bankrupt.[1] What is your intuition concerning the beetle population, assuming there is a positive probability that a female has no offspring in a breeding season? (Hint: as John Maynard Keynes said, "in the long run we are all dead.")

It turns out that there is a very simple and powerful theorem that tells us exactly when a Markov chain is ergodic, and provides a simple characterization of the long-run behavior of the system. To develop the machinery needed to express and understand this theorem, we define a few terms. Let $p_{ij}^{(m)}$ be the probability of being in state j in m periods if the chain is currently in state i. Thus if we start in state i at period 1, the probability of being in state j at period 2 is just $p_{ij}^{(1)} = p_{ij}$. To be in state j in period 3 starting from state i in period 1, the system must move from state i to some state k in period 2, and then from k to j in period 3. This happens with probability $p_{ik}p_{kj}$. Adding up over all k, the probability of being in state j in period 3 is

$$p_{ij}^{(2)} = \sum_k p_{ik}p_{kj},$$

Using matrix notation, this means the matrix of two-period transitions is given by

$$P^{(2)} = \{p_{ij}^{(2)} | i, j = 1, 2, \ldots\} = P^2.$$

Generalizing, we see that the k-period transition matrix is simply P^k. What we are looking for, then, is the limit of P^k as $k \to \infty$. Let us

[1]Specifically, you can show that the probability player 1 wins is $k_1/(k_1 + k_2)$, and if player 1 has wealth w at some point in the game, the probability he will win is $w/(k_1 + k_2)$.

call this limit (supposing it exists) $P* = \{p*_{ij}\}$. Now $P*$ must have two properties. First, since the long-run behavior of the system cannot depend on where it started, $p*_{ij} = p*_{i'j}$ for any two states i and i'. This means that all the rows of $P*$ must be the same. Let us denote the (common value of the) rows by $u = \{u_1, \ldots, u_n\}$, so u_j is the probability the Markov chain will be in state j in the long run. The second fact is that

$$PP* = P \lim_{k\to\infty} P^k = \lim_{k\to\infty} P^{k+1} = P*.$$

This means u must satisfy

$$u_j = \lim_{m\to\infty} p_{ij}^{(m)} \quad \text{for } i = 1, \ldots, n \tag{16.2}$$

$$u_j = \sum_i u_i p_{ij} \tag{16.3}$$

$$\sum_k u_k = 1, \tag{16.4}$$

for $j = 1, \ldots, n$. Note that (16.3) can be written in matrix notation as $u = uP$, so u is a *left eigenvector* of P. The first equation says that u_j is the limit probability of being in state j starting from any state, the second says that the probability of being in state j is the probability of moving from some state i to state j, which is $u_i p_{ij}$, summed over all states i, and the final equation says u is a probability distribution over the states of the Markov chain. The *recursion equations* (16.3) are often sufficient to determine u, which we call the *invariant distribution* or *stationary distribution* of the Markov chain.

A few examples are useful to get a feel for the recursion equations (16.3). Consider first the n-state Markov chain called the *random walk on a circle*, in which there are n states, and from any state $t = 2, \ldots, n-1$ the system moves with equal probability to the previous or the next state, from state n it moves with equal probability to state 1 or state $n - 11$, and from state 1 it moves with equal probability to state 2 and to state n. In the long run, it is intuitively clear that the system will be all states with equal probability $1/n$. To derive this from the recursion equations, note that the probability transition matrix for the problem is given by

$$P = \begin{bmatrix} 0 & 1/2 & 0 & \ldots & 0 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & \ldots & 0 & 0 & 0 \\ & & & \vdots & & & \\ 0 & 0 & 0 & \ldots & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 0 & \ldots & 0 & 1/2 & 0 \end{bmatrix}.$$

The recursion equations for this system are given by

$$\begin{aligned} u_1 &= \frac{1}{2}u_n + \frac{1}{2}u_2 \\ u_i &= \frac{1}{2}u_{i-1} + \frac{1}{2}u_{i+1} \quad i = 2, \ldots, n-1 \\ u_n &= \frac{1}{2}u_1 + \frac{1}{2}u_{n-1} \\ \sum_{i=1}^{n} u_i &= 1. \end{aligned}$$

Clearly this set of equations has solution $u_i = 1/n$ for $i = 1, \ldots, n$ (it can be shown that there are no other solutions, but that is harder).

Consider next a closely related n-state Markov chain called the *random walk on the line with reflecting barriers*, in which from any state $2, \ldots, n-1$ the system moves with equal probability to the previous or the next state, but from state 1 it moves to state 2 with probability one, and from state n it moves to state $n-1$ with probability one. Intuition in this case is a bit more complicated, since states 1 and n behave differently from the other states. The probability transition matrix for the problem is given by

$$P = \begin{bmatrix} 0 & 1 & 0 & \ldots & 0 & 0 & 0 \\ 1/2 & 0 & 1/2 & \ldots & 0 & 0 & 0 \\ & & & \vdots & & & \\ 0 & 0 & 0 & \ldots & 1/2 & 0 & 1/2 \\ 0 & 0 & 0 & \ldots & 0 & 1 & 0 \end{bmatrix}.$$

The recursion equations for this system are given by

$$\begin{aligned} u_1 &= u_2 \\ u_i &= \frac{1}{2}u_{i-1} + \frac{1}{2}u_{i+1} \quad i = 2, \ldots, n-1 \\ u_n &= u_{n-1} \\ \sum_{i=1}^{n} u_i &= 1. \end{aligned}$$

Once again this set of equations has solution $u_i = 1/n$ for $i = 1, \ldots, n$, which represents the stationary distribution for this Markov chain.

For an example of a denumerable Markov chain, suppose a *serial monogamist* is in state k if he is in a marriage that will dissolve in k days, and is in state 0 if he is currently mateless (yes, we can label the states starting

with zero if we please). Suppose when mateless, he remains mateless with probability f_0, and finds a mate to whom he will be married for k days with probability f_k, $k = 1, \ldots, \infty$. This is a Markov chain with $p_{0k} = f_k$ for all k, and $p_{k,k-1} = 1$ for $k \geq 1$, all other transition probabilities being zero. The recursion equations in this case are

$$u_i = u_0 f_i + u_{i+1}$$

for $i \geq 0$. If we let $r_k = f_k + f_{k+1} + \ldots$ (so r_k is the probability of entering a marriage that lasts at least k days), it is easy to see that $u_k = r_k u_0$ satisfies the recursion equations, and the requirement that $\sum_i u_i = 1$ becomes $u_0 = 1/\mu$, where $\mu = \sum r_k$. It is straightforward to show that μ is the expected value of the probability distribution $\{f_k\}$ (show it!), so we conclude that if this expected value does not exist, then no stationary distribution exists. Otherwise the stationary distribution is given by

$$u_i = r_i/\mu \quad i = 0, 1, \ldots .$$

Note that $\mu = 1/u_0$ is the expected number of periods between visits to u_0, since μ is the expected value of $\{f_k\}$. We can also show that $1/u_k = \mu/r_k$ is the expected number of periods μ_k between visits to state k, for any $k \geq 0$. Indeed, the fact that $u_k = 1/\mu_k$, where u_k is the probability of being in state k in the long run and μ_k is the expected number of periods between visits to state k, is a general feature of Markov chains with stationary distributions, as we show below.

Let us prove that $\mu_k = \mu/r_k$ for $k = 2$ in the above model, leaving the general case to the reader. From state 2 the Markov chain moves to state 0 in two periods, then requires some number j of periods before it moves to some state $k \geq 2$, and then in $k - 2$ states moves to state 2. Thus if we let v be the expected value of j and we let w represent the expected value of k, we have $\mu_k = 2 + v + w - 2 = v + w$. Now v obviously satisfies the recursion equation

$$v = f_0(1 + v) + f_1(2 + v) + r_2(1),$$

since after a single move the system remains in state 0 with probability f_0 and the expected number of periods before hitting $k > 1$ is $1 + v$ (the first term), or moves to state 1 with probability f_1 and the expected number of periods before hitting $k > 1$ is $2 + v$ (the second term), or hits $k > 1$ immediately with probability r_2 (the final term),. Solving, we find that $v = (1 + f_1)/r_2$. To find w, note that the probability of being

in state k conditional on $k \geq 2$ is f_k/r_2. Thus

$$\begin{aligned} w &= \frac{1}{r_2}(2f_2 + 3f_3 + \ldots) \\ &= \frac{1}{r_2}(\mu - f_0 - 2f_1 - r_2) \\ &= \frac{1}{r_2}(\mu - 1 - f1). \end{aligned}$$

Thus $v + w = \mu/r_2$.

16.19.1 The Ergodic Theorem for Markov Chains

When are the equation (16.2) true, and what exactly do they say? To answer this, we will need a few more concepts. We limit our discussion to *irreducible Markov chains*, which we define as follows. We say a state j can be *reached* from a state i if $p_{ij}^{(m)} > 0$ for some positive integer m. We say a Markov chain is *irreducible* if every state can be reached from every other state. If a Markov chain M is not irreducible, we can carve up M into sets of states each of which represents an irreducible Markov chain, and states in each subset can be reached only by states in the same subset (work this out!). Clearly a non-irreducible chain cannot be ergodic, because where it starts out determines the set of states in which it remains forever.

If M is irreducible, then all the u_i in (16.2) are *strictly positive*. To see this suppose some $u_j = 0$. Then by (16.3), if $p_{ij} > 0$, then $p_i = 0$ Thus any state that reaches j in one period must also have weight zero in u. But a state i' that reaches j in two periods must pass through a state i that reaches j in one period, and since $u_i = 0$, we also must have $u_{i'} = 0$ Extending this argument, we that any state i that reaches j must have $u_i = 0$, and since M is irreducible, all the $u_i = 0$, which violates (16.4). The fact that the u_i must be strictly positive is central to the following argument.

We say a Markov chain is *periodic* if there is some integer $k > 1$ such that $p_{ii}^{(k)} > 0$ and $p_{ii}^{(m)} = 0$ unless m is a multiple of k; i.e., the Markov chain can return to state i only after time periods that are multiples of some integer greater than one. If a state is not periodic, we say it is *aperiodic*. Clearly in an irreducible Markov chain M, all states must be aperiodic or the limit in (16.2) could not exist—some subsequences of $p_{jj}^{(m)}$ would have limit zero, and some would have a positive limit, so u_j could not satisfy (16.2).

Now let q_i be the probability that, starting from state i, the system returns to state i in some future period. If $q_i < 1$, then it is clear that state i can only occur a finite number of times, so in the long run $u_i = 0$, which we have seen is impossible for an irreducible Markov chain. We say a state i is *persistent* if $q_i = 1$, and *transient* otherwise.

Aperiodicity and persistence are enough for finite Markov chains, but for denumerable Markov chains another complication can arise: a state i can be persistent, but the expected number of periods μ_i before a return to state i could be infinite. We call μ_i the *mean recurrence time* of state i. If the mean recurrence time of state i is μ_i, M is in state i on average one period out of every μ_i, so we should have $u_i = 1/\mu_i > 0$. We found this to be true in the case of the serial monogamist above, and in fact it can be shown to be true in general. Then M can have a stationary distribution only if μ_i is finite. This motivates the following definition: an aperiodic persistent state i is called *ergodic* if the mean recurrence time μ_i is finite.

This argument motivates the following *ergodic theorem for Markov chains*, the proof of which can be found in Feller (1950):

Theorem 16.1. *Let M be an irreducible Markov chain with no periodic states. If all states in M are ergodic, then the recursion equations (16.2) hold, and all $u_i > 0$. Conversely, suppose $u_1, \ldots, u_n \geq 0$ exist satisfying the recursion equations (16.3) and (16.4). Then all states are ergodic, (16.2) holds, and $u_j = 1/\mu_j > 0$ for each state j, where μ_j is the mean recurrence time for state j.*

16.19.2 The Sisyphusian Markov Chain

As an exercise, consider the following *Sisyphusian Markov chain*, in which Albert has a piano on his back and must climb up and infinite number of steps $k = 1, 2, \ldots$. At step k, with probability b_k he stumbles and falls all the way back to the first step, and with probability $1 - b_k$ he proceeds to the next step. This gives the probability transition matrix

$$P = \begin{bmatrix} b_1 & 1-b_1 & 0 & 0 & 0 & \ldots \\ b_2 & 0 & 1-b_2 & 0 & 0 & \ldots \\ b_3 & 0 & 0 & 1-b_3 & 0 & \ldots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}.$$

(a) Show that the recursion equations for this system are

$$u_1 = \sum_i u_i b_i$$

$$u_{k+1} = u_k(1 - b_k).$$

(b) Show that these equations are satisfied only if

$$b_1 + (1 - b_1)b_2 + (1 - b_1)(1 - b_2)b_3 + \ldots = 1,$$

which is true unless the left hand side is infinite. Hint: Let the left hand side equal $x < \infty$, subtract b_1 from both sides and divide by $1 - b_1$; now the left hand side does not depend on b_1 but the right hand side does, unless $x = 1$.

(c) Use the fact that $\sum_i u_i = 1$ to find u_1, which must satisfy

$$u_1[1 + (1 - b_1)b_2 + (1 - b_1)(1 - b_2)b_3 + \ldots] = 1.$$

(d) Show that the Markov chain is ergodic if $b_k = 1/2$ and find the $\{u_i\}$.

(e) Show that the Markov chain is not ergodic if $b_k = 1/k$, because the mean time between passages to state 1 is infinite.

(f) What are the implications of this analysis for existentialist philosophy?

16.19.3 Andrei Andreyevich's Two Urn Problem

After Andrei Andreyevich Markov discovered the chains for which he is named, he proved the Ergodic Theorem (see above) for finite chains. Then he looked around for an interesting problem to solve. Here is what he came up with (this problem had been solved before, but not rigorously).

Suppose there are two urns, one teal and one magenta. Each containing m balls. Of the $2m$ balls, b are khaki and the remainder are olive. At each time period $t = 1, \ldots$ two balls are drawn randomly, one from each urn, and each ball is placed in the other urn. Let state i represent the event that there are $i \in [0, \ldots, b]$ khaki balls in the teal urn. What is the probability u_i of state i in the long run?

Let $P = \{p_{ij}\}$ be the $(b + 1) \times (b + 1)$ probability transition matrix. To move from i to $i - 1$, a khaki ball must be drawn from the teal urn, and an olive ball must be drawn from the magenta urn. This means $p_{i,i-1} = i(m - b + i)/m^2$. To remain in state i, either both balls drawn are khaki or both are teal, $p_{i,i} = (i(b - i) + (m - i)(m - b + i)/m^2$. To

move from i to $i+1$, an olive ball must be drawn from the teal urn, and a khaki ball must be drawn from the magenta urn. This means $p_{i,i+1} = (m-i)(b-i)/m^2$. All other transition probabilities are zero.

The recursion equations in this case are given by

$$u_i = u_{i-1}p_{i-1,i} + u_i p_{ii} + u_{i+1}p_{i+1,i}.$$

for $i = 0, \ldots, b+1$, where we set $u_{-1} = u_{b+2} = 0$. I do not know how Andrei solved these equations, but I suspect he guessed at the answer and then showed that it works. At any rate that is what I shall do. Our intuition concerning the Ergodic Theorem suggests that in the long run the probability distribution of khaki balls in the teal urn are the same as if m balls were randomly picked from a pile of $2m$ balls (of which b are khaki) and put in the teal urn. If we write the number of combinations of n things taken r at a time as $\binom{n}{r} = n!/r!(n-r)!$, then u should satisfy

$$u_i = \binom{m}{i}\binom{m}{b-i}/\binom{2m}{b}.$$

The denominator in this expression is the number of ways the b khaki balls can be allocated to the $2m$ possible positions in the two urns, and the numerator is the number of ways this can be done when i khaki balls are in the teal urn. You can check that u now satisfies the recursion equations.

16.20 Preferences and Expected Utility

Game theory is based on the notion that agents maximize utility subject to constraints (e.g., in choosing a best response in a game), but no one has actually ever discovered a utility function, and most behavioral scientists do not believe they exist. Rather, the claim is that if the behavior of an agent has certain regularities, we can fully describe the agent's behavior *as if* it were the result of maximizing utility.[2] In this section we develop a set of behavioral properties (we will call them 'axioms') that yield not only a utility function over outcomes, but a probability distribution over states of nature, and such that the expected utility principle holds. We outline Savage's (1954) classic analysis of this problem.[3]

[2]This quite nicely parallels physics, where the behavior of a system can be described by finding the maxima or minima of a function (e.g., total energy, or the length of a path), with no implication that this actually describes the physical processes involved.

[3]For a more leisurely but still relatively informal presentation of this material, see Kreps 1988. For a complete treatment, see Savage 1954 or Fishburn 1970.

Let X be a set of payoffs, with characteristic elements x, y, x', y'. For ease of exposition, we shall assume X is a subset of some $\mathbf{R}^n$, though this is not strictly necessary. Suppose when faced with the choice between x and y, the agent is observed to choose x. We say x is *weakly preferred* to y, and we write $x \succeq y$. We say 'weakly' because we cannot tell simply by observing that x was chosen whether the agent strictly prefers x to y, or is in fact indifferent between the two, and on another occasion might choose y over x under the same conditions. We call $\succeq$ a *binary relation*.[4] Indeed, we will assume $\succeq$ is *reflexive*, meaning that for all $x \in X$, $x \succeq x$.

Given $\succeq$, we can define a bunch of other binary relations, including the following: (a) $x \not\succeq y$ means "not $x \succeq y$"; (b) $x \succ y$ means "$x \succeq y$ and not $y \succeq x$," and we say "x is strictly preferred to y"; (c) $x \not\succ y$ means "not $x \succ y$"; (d) $x \sim y$ means "$x \succeq y$ and $y \succeq x$," and we say "the agent is indifferent between x and y"; (e) $x \not\sim y$ means "not $x \sim y$." These definitions introduce no new concepts, but make it easier to talk about preferences relations.

As an exercise, check for yourself the following: (a) for any $x, y \in X$, exactly one of the relations $x \succ y$, $x \sim y$, and $y \succ x$ holds; (b) $x \succeq y$ is equivalent to $x \succ y$ or $x \sim y$; (c) $\sim$ and $\succ$ are transitive, but neither is complete; (d) $\succ$ is *asymmetric*, meaning $x \succ y$ implies $y \not\succ x$; (e) $\sim$ is *symmetric*, meaning $x \sim y$ implies $y \sim x$.

We say $u : X \to \mathbf{R}$ is a *utility function* representing $\succeq$ if, for any $x, y \in X$, $x \succ y$ if and only if $u(x) > u(y)$. It is clear that if $\succeq$ is represented by a utility function $u(\cdot)$, then for any $x, y \in X$, $x \succeq y$ if and only if $u(x) \geq u(y)$. Also for any $x, y \in X$, we must have either $x \succeq y$ or $y \succeq x$, since either $u(x) \geq u(y)$ or $u(y) \geq u(x)$. If $\succeq$ has this property, we say that $\succeq$ is a *complete binary relation*. Moreover, if $x, y, z \in X$, $x \succeq y$ and $y \succeq z$, then $x \succeq z$, since if $u(x) \geq u(y)$ and $u(y) \geq u(z)$ then $u(x) \geq u(z)$. If $\succeq$ has this property, we say that $\succeq$ is a *transitive binary relation*. We say $\succeq$ is a *preference relation* if it is both complete and transitive. For the rest of this section we assume $\succeq$ is a preference relation. To insure that the analysis is not trivial, we also assume that $x \succeq y$ is false for at least some $x, y \in X$. It is easy to show that

Theorem 16.2. *A binary relation $\succeq$ on the set X of payoffs can be represented by a utility function $u : X \to \mathbf{R}$ if and only if $\succeq$ is a preference relation.*

[4]More generally, a binary relation on X is an operator $\odot$ such that for any $x, y \in X$, $x \odot y$ is either true or false. For instance '$>$' is a binary relation in arithmetic, as are $<, \leq, =, \geq$, and $\neq$. But '$+$' is not a binary relation, since $x + y$ is neither true nor false.

This shows that we can use utility functions to represent choice behavior whatever the underlying mechanism driving choice, so long as choice is complete (i.e., the agent can always choose between two payoffs) and transitive (i.e., the agent never weakly prefers x to y, weakly prefers y to z, but strongly prefers z to x). It is clear that $u(\cdot)$ is not unique, and indeed, we have

Theorem 16.3. *If $u(\cdot)$ represents the preference relation $\succeq$ and $f(\cdot)$ is a strictly increasing function, then $v(\cdot) = f(u(\cdot))$ also represents $\succeq$. Conversely if both $u(\cdot)$ and $v(\cdot)$ represent $\succeq$ then there is an increasing function $f(\cdot)$ such that $v(\cdot) = f(u(\cdot))$.*

The first half of the theorem is true because if f is strictly increasing, then $u(x) > u(y)$ implies $v(x) = f(u(x)) > f(u(y)) = v(y)$, and conversely. For the second half, suppose $u(\cdot)$ and $v(\cdot)$ both represent $\succeq$, and for any $y \in \mathbf{R}$ such that $v(x) = y$ for some $x \in X$, let $f(y) = u(v^{-1}(y))$. Then $f(\cdot)$ is increasing (since it is the composition of two increasing functions) and $f(v(x)) = u(v^{-1}(v(x))) = u(x)$, which proves the theorem. Q.E.D.

Savage's accomplishment was to show that if the agent has a preference relation over *lotteries* that has some plausible properties, then not only can the agent's preferences be represented by a utility function, but we can infer the probabilities the agent implicitly places on various events, and the expected utility principle holds for these probabilities.

To see this, let S be a finite set of *states of nature*. We call subsets $A \subseteq S$ *events*. Also, let $\mathcal{L}$ be a set of 'lotteries,' where a *lottery* is a function $\pi : S \to X$ that associates with each state of nature $s \in S$ a payoff $\pi(s) \in X$. We suppose that the agent chooses among lotteries without knowing the state of nature, after which 'Nature' choses the state $s \in S$ that obtains, so that if the agent chose lottery $\pi \in \mathcal{L}$, his payoff is $\pi(s)$.

Now suppose the agent has a preference relation $\succ$ over $\mathcal{L}$ (we use the same symbol $\succ$ for preferences over both outcomes and lotteries). We seek a set of plausible properties of $\succ$ over lotteries that together allow us to deduce (a) a utility function $u : X \to \mathbf{R}$ corresponding to the preference relation $\succ$ over outcomes in X; (b) a probability distribution $p : S \to \mathbf{R}$ such that the expected utility principle holds with respect to the preference relation $\succ$ over lotteries and the utility function $u(\cdot)$; i.e., if we define

$$\mathbf{E}_\pi[u; p] = \int_S p(s)u(\pi(s)), \tag{16.5}$$

where the integral is suitably defined over S, then for any $\pi, \rho \in \mathcal{L}$

$$\pi \succ \rho \qquad \Leftrightarrow \qquad \mathbf{E}_\pi[u; p] > \mathbf{E}_\rho[u; p].$$

Our first condition is that $\pi \succ \rho$ depends only on states of nature where π and ρ have different outcomes. We state this more formally as

(**A1**) For any $\pi, \rho, \pi', \rho' \in \mathcal{L}$ and any event $A \subseteq S$, if $\pi(s) = \rho(s)$ and $\pi'(s) = \rho'(s)$ for any $s \in A$, and $\pi(s) = \pi'(s)$ and $\rho(s) = \rho'(s)$ for any $s \in A^c$, then $\pi \succ \rho \Leftrightarrow \pi' \succ \rho'$

(recall that $A^c = \{s \in S | s \notin A\}$). This axiom allows us to define a *conditional preference* $\pi \succ_A \rho$, where $A \subseteq S$, which we interpret as "π is strictly preferred to ρ, given event A," as follows. We say $\pi \succ_A \rho$ if, for some $\pi', \rho' \in \mathcal{L}$, $\pi(s) = \pi'(s)$ and $\rho(s) = \rho'(s)$ for all $s \in A$, and $\pi'(s) = \rho'(s)$ for $s \notin A$. This allows us to define $\succeq_A$ and $\sim_A$ in a similar manner. We then define an event $A \subseteq S$ to be *null* if $\pi \sim_A \rho$ for all $\pi, \rho \in \mathcal{L}$.

Our second condition is then

(**A2**) If $A \subseteq S$ is not null, and if $\pi(s) = x$ and $\rho(s) = y$ for all $s \in A$, then $\pi \succ_A \rho \Leftrightarrow x \succ y$.

This axiom says that a natural relationship between outcomes and lotteries holds: if π pays x given event A and ρ pays y given event A, and if $x \succ y$, then $\pi \succ_A \rho$, and conversely.

Our third condition asserts that the probability that a state of nature occurs is independent from the outcome one receives when the state occurs. The difficulty in stating this axiom is that the agent cannot choose probabilities, but only lotteries. But if the agent prefers x to y, and if $A, B \subseteq S$ are events and $x \succ y$, then the agent treats A as 'more probable' than B if and only if a lottery that pays x when A occurs and y when A does not occur will be preferred to a lottery that pays x when B occurs and y when B does not. However this must be true for any $x, y \in X$ such that $x \succ y$, or the agent's notion of probability is incoherent (i.e., it depends on what particular payoffs we are talking about—for instance, 'wishful thinking,' where if the prize associate with an event increases, the agent thinks it is more likely to occur). More formally, we have

(**A3**) Suppose $x \succ y$, $x' \succ y'$, $\pi, \rho, \pi', \rho' \in \mathcal{L}$, and $A, B \subseteq S$. Suppose that π pays off x when $s \in A$ and y when $s \in A^c$, and ρ pays off x' when $s \in A$ and y' when $s \in A^c$. Moreover,

suppose that π' pays off x when $s \in B$ and y when $s \in B^c$, and ρ' pays off x' when $s \in B$ and y' when $s \in B^c$. Then $\pi \succ \rho \Leftrightarrow \pi' \succ \rho'$.

The fourth condition is a weak version of the notion that if one lottery has a higher payoff than another for any event, then the first is preferred to the second. To state this condition, note that if $x \in X$, we can also consider $x \in \mathcal{L}$ to be the lottery that pays x for all states of nature. We then have

(**A4**) For any event A, if $\pi \succ_A \rho(s)$ for all $s \in A$, then $\pi \succeq_A \rho$. Also, any event A, if $\rho(s) \succ_A \pi$ for all $s \in A$, then $\rho \succeq_A \pi$.

In other words, if for any event A, π pays more than the best ρ can pay, the $\pi \succ \rho$, and conversely.

Finally, we need a technical property that is something like the density property that is needed to show that a preference relation can be represented by a utility function. It says that for any $\pi, \rho \in \mathcal{L}$, and any $x \in X$, we can partition S into a number of disjoint subsets $A_1, \ldots A_n$ such that $\cup_i A_i = S$, and for each A_i, if we change π so that its payoff is x on A_i, then π is still preferred to ρ, and similarly, for each A_i, if we change ρ so that its payoff is x on A_i, then π is still preferred to ρ. This means that no payoff is 'supergood,' so that no matter how unlikely an event A is, a lottery with that payoff is always preferred to a lottery with a different payoff. Similarly, no payoff can be 'superbad.' The condition is formally as follows:

(**A5**) For all $\pi, \rho \in \mathcal{L}$ with $\pi \succ \rho$, and for all $x \in X$, there is are disjoint subsets $A_1, \ldots, A_n$ of S such that $\cup_i A_i = S$ and for any A_i (a) if $\pi'(s) = x$ for $s \in A_i$ and $\pi'(s) = \pi(s)$ for $s \in A^c$, then $\pi' \succ \rho$, and (b) if $\rho'(s) = x$ for $s \in A_i$ and $\rho'(s) = \rho(s)$ for $s \in A^c$, then $\pi \succ \rho'$.

We then have Savage's theorem:

Theorem 16.4. *Suppose* **A1-A5** *hold. Then there is a probability function p on S and a utility function $u : X \to \mathbf{R}$ such that for any $\pi, \rho \in \mathcal{L}$, $\pi \succ \rho$ if and only if $\mathbf{E}_\pi[u; p] > \mathbf{E}_\rho[u; p]$.*

The proof of this theorem is somewhat arduous (it is sketched in Kreps, 1988).

16.21 Exceptions to the Expected Utility Principle

Although most economists consider the expected utility principle acceptably accurate as a basis of modeling behavior—it is, after all, quite central to all of game theory—there are certainly well-established situations in which agents appear to violate it. Machina (1987) reviews this body of evidence and presents models to deal with them We sketch here the most famous of these 'anomalies,' the so-called *Allais paradox* and *Ellsberg paradox*. They are of course not 'paradoxes' at all, but simply empirical regularities that do not fit the expected utility principle.

Maurice Allais (1953) offered the following scenario. There are two choice situations in a game with prizes (\$2,500,000, \$500,000, \$0). The first is a choice between a lottery π that offers these prizes with probabilities (0,1,0) and π' that offers the prizes with probabilities (0.10,0.89,0.01). The second is a choice between a lottery ρ with probabilities (0,0.11,0.89) and ρ' with probabilities (0.10,0,0.90). Most people, when faced with these two choice situations choose $\pi \succ \pi'$ and $\rho' \succ \rho$ (what would you choose?). It is easy to see that this is not consistent with the expected utility principle. To see this, let us write $u_h = u(2,500,000)$, $u_m = u(500,000)$, and $u_l = u(0)$. Then if the expected utility principle holds, $\pi \succ \pi'$ implies $u_m > 0.1u_h + 0.89u_m + 0.01u_l$, so $0.11u_m > 0.10u_h + 0.01u_l$, which implies (adding $0.89u_l$ to both sides) $0.11u_m + 0.89u_l > 0.10u_h + 0.9u_l$, which says $\rho \succ \rho'$.

Why do people make this 'mistake?' Perhaps because of *regret*, which does not mesh well with the expected utility principle (Loomes 1988, Sugden 1993). If you choose π' in the first case and you end up getting nothing, you will feel really foolish, whereas in the second case, you are probably going to get nothing anyway (not your fault), so increasing the chances of getting nothing a tiny bit (0.01) gives you a good chance (0.10) of winning the really big prize. Or perhaps *loss aversion* (§11.3.3), since in first case, the anchor point (the most likely outcome) is \$500,000, while in the second case the anchor is 0. Loss averse agents will then shun π', which give a positive probability of loss, whereas in the second case, neither lottery involves a loss, from the point of view of the most likely outcome.

Another classic violation of the expected utility principle was suggested by Ellsberg (1961). Consider two urns. Urn A has 51 red balls and 49 white balls. Urn B also has 100 red and white balls, but the fraction of red balls is unknown. One ball is chosen from each urn, but remains hidden from sight. Subjects are asked to choose in two situations. First,

a subject can choose the ball from urn A or urn B, and if the ball is red, he wins \$10. In the second situation, the subject can choose the ball from urn A or urn B, and if the ball is white, he wins \$10. Many subjects choose the ball from urn A in both cases. This violates the expected utility principle, no matter what probability the subject places on the probability p that the ball from urn B is white. For in the first situation the payoff to choosing urn A is $0.51u(10) + 0.49u(0)$ and the payoff from choosing urn B is $(1-p)u(10)+pu(0)$, so strictly preferring urn A means $p < 0.49$. In the second situation the payoff to choosing urn A is $0.49u(10) + 0.51u(0)$ and the payoff from choosing urn B is $pu(10) + (1-p)u(0)$, so strictly preferring urn A means $p > 0.49$. This shows that the expected utility principle does not hold.

The usual explanation of this behavior is that the subject *knows* the probabilities associated with the first urn, while the probabilities associated with the second urn are *unknown*. People prefer the known to the unknown, then, in a way that violates the expected utility principle.

16.22 Risk Behavior and the Shape of the Utility Function

If $\succeq$ is defined over X, we can say nothing about the *shape* of a utility function $u(\cdot)$ representing $\succeq$, since by Theorem 16.3, any increasing function of $u(\cdot)$ also represents $\succeq$. However if $\succeq$ is represented by a utility function $u(x)$ satisfying the expected utility principle, then $u(\cdot)$ is determined up to an arbitrary constant and unit of measure.[5] Specifically, we have

Theorem 16.5. *Suppose the utility function $u(\cdot)$ represents the preference relation $\succeq$ and satisfies the expected utility principle. If $v(\cdot)$ is another utility function representing $\succeq$, then there are constants $a, b \in \mathbf{R}$ with $a > 0$ such that $v(x) = au(x) + b$ for all $x \in X$.*

For a proof of this theorem, see Mas-Colell, Whinston, and Green (1995), p. 173.

Supposing $X = \mathbf{R}$, so the payoffs can be considered to be 'money,' and utility satisfies the expected utility principle, what shape do such utility functions have? It would be nice if they were linear in money, in which case expected utility and expected value would be the same thing (why?). But the famous Swiss mathematician Daniel Bernoulli (1700–1782), who participated in the birth of probability theory, showed that

[5] According to long-standing tradition that has little foundation in mathematics, we say utilities over outcomes are *ordinal*, while utilities over lotteries are *cardinal*.

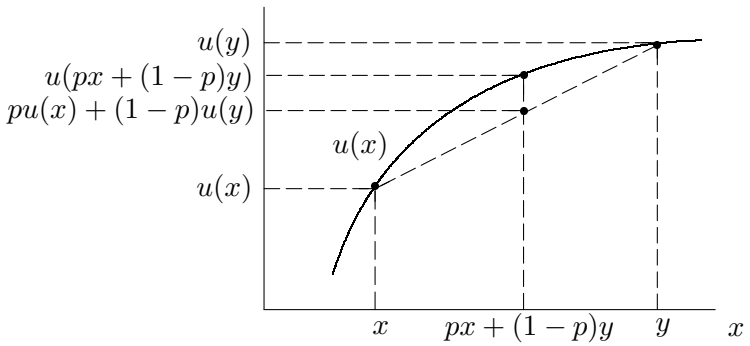

Figure 16.1: A Concave Utility Function

linear utility is highly implausible in his famous *St. Petersberg Paradox*. Consider a lottery in which I toss a coin repeatedly times until it comes up tails. If the total number of heads is k, then you win 2^k dollars. How much would you pay me to play this lottery? Notice that if the coin comes up tails on the first toss, you get nothing. But if I toss ten heads in a row before a tail, you get $2^{10} = 1024$ dollars. I have found that most students are willing to pay between \$2.00 and \$10.00 to play this game. But, as Bernoulli showed, the expected value of the lottery is *infinite*, since the payoff is

$$\frac{1}{2} \cdot 2 + \frac{1}{4} \cdot 4 + \frac{1}{8} \cdot 8 + \ldots = 1 + 1 + 1 + \ldots = \infty.$$

This example shows that utility must grow less fast than a linear function, which means it is probably concave, as illustrated in Figure 16.1. We say a function $u : X \to \mathbf{R}$ is *strictly concave* if, for any $x, y \in X$, and any $p \in (0, 1)$, we have $pu(x) + (1-p)u(y) < u(px + (1-p)y)$. We say $u(x)$ is *weakly concave*, or simply *concave*, if $u(x)$ is either strictly concave or linear, in which case the above inequality is replaced by $pu(x) + (1-p)u(y) = u(px+(1-p)y)$. Notice that if we define the lottery π as paying x with probability p and y with probability $1-p$, then the condition for strict concavity says that *the expected utility of the lottery is less than the utility of the expected value of the lottery*, as depicted in Figure 16.1.

What *are* good candidates for $u(x)$? It is easy to see that strict concavity means $u''(x) < 0$, providing $u(x)$ is twice differentiable (which we assume). But there are lots of functions with this property. According to the famous *Weber-Fechner Law* of psychophysics, for a wide range of sensory stimuli, and over a wide range of levels of stimulation, a just-noticeable change in a stimulus is a constant fraction of the original

stimulus. If this holds for money, then the utility function is logarithmic. Using a convenient scaling (cf. Theorem 16.5), we can write $u(x) = \ln(x)$, so $u(1) = 0$ and $u(0) = -\infty$. This utility function is certainly strictly concave, since $u''(x) = -1/x^2 < 0$ for $x > 0$. The expected utility of Daniel Bernoulli's lottery for this utility function is then

$$\sum_{k=1}^{\infty} \frac{\ln(2^k)}{2^k} = \ln(2),$$

so a person with this utility function would only be willing to pay \$2 to play this game![6] Another plausible concave utility function is $u(x) = x^a$ where $a \in (0, 1)$. This is increasing and $u''(x) = a(a-1)x^{a-2} < 0$, so it is concave. In this case the Bernoulli lottery is

$$\sum_{k=1}^{\infty} \frac{2^{ak}}{2^k} = \frac{1}{1 - 2^{a-1}} - 1,$$

so for instance if $a = 1/2$, so $u(x) = \sqrt{x}$, the utility of the lottery is $\sqrt{2} + 1$, and the amount you would pay to play the lottery would be the square of this number, which is about \$5.83.

We say an agent is *risk averse* if the agent prefers the expected value of a lottery to the lottery itself (provided, of course, the lottery does not offer a single payoff with probability one, which we call a 'sure thing'). We know, then, that an agent with utility function $u(\cdot)$ is risk averse if and only if $u(\cdot)$ is concave.[7] Similarly, we say an agent is *risk loving* if he prefers any lottery to the expected value of the lottery, and *risk neutral* if he is indifferent between a lottery and its expected value. Clearly an agent is risk neutral if and only if he has linear utility.

Does there exist a measure of risk aversion that allows us to say when one individual is more risk averse than another, or how an agent's risk aversion changes with changing wealth? We may define agent A to be

[6] Where did the ln(2) answer come from? You can take my word for it, or write out the first ten terms and see what you get. Or, if you like math, you can replace the 2 by x in the sum, integrate the series, find the sum, differentiate the series, and let $x = 2$.

[7] It may be asked why people play government-sponsored lotteries, or spend money at gambling casinos, if people are generally risk averse. The most plausible explanation is that people enjoy the act of gambling. The same woman who will have insurance on her home and car, both of which presume risk aversion (§3.16), will gamble small amounts of money for recreation. An excessive love for gambling, of course, leads an individual either to personal destruction or wealth and fame (usually the former).

more risk averse than agent B if whenever A prefers a lottery to an amount of money x, B will also prefer the lottery to x. We say A is *strictly* more risk averse than B if he is more risk averse, and there is some lottery that B prefers to an amount of money x, but such that A prefers x to the lottery.

Clearly the degree of risk aversion depends on the curvature of the utility function (by definition the *curvature* of $u(x)$ at x is $u''(x)$), but since $u(x)$ and $v(x) = au(x) + b$ $(a > 0)$ describe the same behavior, but $v(x)$ has curvature a time that of $u(x)$, we need something more sophisticated. The obvious candidate is $\lambda_u(x) = -u''(x)/u'(x)$, which does not depend on scaling factors. This is called the *Arrow-Pratt coefficient of absolute risk aversion*, and it is exactly the measure that we need. We have

Theorem 16.6. *An agent with utility function $u(x)$ is more risk averse than an agent with utility function $v(x)$ if and only if $\lambda_u(x) > \lambda_v(x)$ for all x.*

For example, the logarithmic utility function $u(x) = \ln(x)$ has Arrow-Pratt measure $\lambda_u(x) = 1/x$, which decreases with x; i.e., as the agent becomes wealthier, he becomes less risk averse. Studies show that this property, called *decreasing absolute risk aversion*, holds rather widely (Rosenzweig and Wolpin 1993, Saha, Shumway, and Talpaz 1994, Nerlove and Soedjiana 1996). Another increasing concave function is $u(x) = x^a$ for $a \in (0, 1)$, for which $\lambda_u(x) = (1 - a)/x$, which also exhibits declining absolute risk aversion. Similarly, $u(x) = 1 - x^{-a}$ $(a > 0)$ is increasing and concave, with $\lambda_u(x) = -(a + 1)/x$, which again exhibits declining absolute risk aversion This utility has the additional attractive property that *utility is bounded*: no matter how rich you are, $u(x) < 1$.[8] Yet another candidate for a utility function is $u(x) = 1 - e^{-ax}$, for some $a > 0$. In this case $u(x) = a$, which we call *constant absolute risk aversion*.

Another commonly used term is the *coefficient of relative risk aversion*, $\mu_u(x) = \lambda_u(x)/x$. Notice that for any of the utility functions $u(x) = \ln(x)$, $u(x) = x^a$ for $a \in (0, 1)$, and $u(x) = 1 - x^{-a}$ $(a > 0)$, $\mu_u(x)$ is constant, which we call *constant relative risk aversion*. For $u(x) = 1 - e^{-ax}$ $(a > 0)$, we have $\mu_u(x) = a/x$, so we have *declining relative risk aversion*.

[8] If utility is unbounded, it is easy show that there is a lottery that you would be willing to give all your wealth to play, no matter how rich you are. This is not plausible behavior.

As an exercise, consider the utility function

$$u(x) = 1 - e^{-ax^b},$$

where $b < 1$ and $ab > 0$.

(a) Show that $u(x)$ is increasing and concave;

(b) find the Arrow-Pratt measures of absolute and relative risk aversion for this utility function;

(c) show that $u(x)$ exhibits decreasing absolute risk aversion for $a < 0$, constant relative risk aversion for $a = 0$, and increasing relative risk aversion for $a > 0$.